NAS Technical Report; NAS-06-016
September 2006

Role of High-End Computing in Meeting NASA's Science and Engineering Challenges

Rupak Biswas, Eugene L. Tu, and William R. Van Dalsem

NASA Ames Research Center, Moffett Field, CA 94035, USA
{Rupak.Biswas, Eugene.L.Tu, William.R.VanDalsem}@nasa.gov

Abstract
High-end computing (HEC) has always played a major role in meeting the modeling and simulation needs of various NASA missions. Two years ago, NASA was on the verge of dramatically enhancing its HEC capability and capacity by significantly increasing its computational and storage resources. With the 10,240-processor Columbia supercomputer in production since October 2004, HEC is having an even greater impact within the Agency and beyond. Advanced science and engineering simulations in space exploration, Shuttle operations, Earth sciences, and fundamental aeronautics research are occurring on Columbia, demonstrating its ability to accelerate NASA's exploration vision. This paper describes how the integrated production environment fostered at the NASA Advanced Supercomputing (NAS) facility at Ames Research Center is reducing design cycle times, accelerating scientific discovery, achieving rapid parametric analyses of multiple scenarios, and enhancing safety for several NASA missions. We focus on Columbia's impact on two key engineering and science disciplines: aerospace, and climate/weather. We also discuss future mission challenges and plans for NASA's next-generation HEC environment.

1 Introduction

Over the years, high-end computing (HEC) has played a major role in meeting the modeling and simulation needs of various NASA missions. Two years ago, having projected its near-term and future high-fidelity computational requirements, NASA was on the verge of dramatically increasing its HEC capability and capacity [3]. With NASA's 10,240-processor supercomputer, Columbia, in production since October 2004, HEC is having an even greater impact within the Agency and extending to partner institutions. Significant cutting-edge science and engineering simulations in the areas of space exploration, Shuttle operations, Earth sciences, and fundamental aeronautics research are occurring regularly on Columbia, demonstrating its ability to accelerate NASA's vision for space exploration [6, 7]. This paper describes how the integrated production environment fostered at the NASA Advanced Supercomputing (NAS) facility located at Ames Research Center is being used to design future aerospace vehicles, conduct parametric analysis for safe operation of the Shuttle, accelerate scientific discovery, and enhance crew safety during the life cycle of NASA missions. Columbia's impact is illustrated using two of the agency's key engineering and science disciplines: aerospace and climate/weather.

In aerospace, computed results are presented in three areas: debris transport analysis for the Space Shuttle Launch Vehicle (SSLV); flowliner analysis for the Space Shuttle Main Engine (SSME); and risk assessment of ascent abort scenarios for proposed Crew Exploration Vehicle (CEV) designs. Among NASA's applications in climate and weather modeling are next-generation global ocean models that resolve eddies and other narrow current systems, and atmospheric modeling and prediction of hurricane tracks for early warning. Columbia is also having a significant impact on NASA's numerous space and exploration applications, such as the development of the Crew Launch Vehicle (CLV), and risk assessment throughout the entire mission cycle—from ground operations, vehicle launch, and return to Earth.

The role of the Columbia supercomputer (currently ranked the fourth fastest system in the world [24], at 62 TFlop/s peak performance) in advancing the science and technologies related to the above topics are illustrated through various data analysis methods. Users of Columbia are also supported by the NAS facility's integrated HEC environment. In addition to system analysts, experts in code parallelization and performance optimization, high-fidelity modeling and simulation, high-speed networking, and data analysis and visualization exploit the power of Columbia to enhance NASA's computational capability.

As with other federal agencies and the commercial sector, NASA's future mission challenges require even more powerful computing systems. At present, the Agency is planning its next-generation HEC system, which will augment Columbia and provide even more high-performance computing resources. Development of this future environment includes anticipated storage and archive requirements for a balanced system, application performance enhancement tools, and faster wide area networking technologies.

2 Aerospace Applications

2.1 Space Shuttle Launch Vehicle Debris Transport Analysis

After the STS-107 incident in February 2003, the Columbia Accident Investigation Board (CAIB) requested that CFD researchers conduct a detailed debris transport analysis to provide insight into the actual mechanism of debris shedding from the bi-pod ramp region of the SSLV. The analysis would also furnish input for foam velocity and density for impact testing to determine damage to the orbiter's wing leading edge [18].

Subsequently, researchers at NASA Ames developed a CFD process for determining the aerodynamic characteristics of debris shedding during the

SSLV ascent phase. In 2005, the role of Columbia's predecessor system in the accident investigation [3] was reported. Since that time, a complete debris scenario has been conducted on Columbia, which focused on predicting the aerodynamic characteristics of potential debris sources, such as insulating foam and ice. This computational analysis was critical to NASA's Return-to-Flight (RTF) effort, where scientists performed six-degree-of-freedom (6-DOF) foam debris transport analyses and visualization to forecast Shuttle damage, and for damage assessment and repair recommendations during the successful Discovery flight in July 2005. HEC and CFD will continue to play a key role for remaining Shuttle flights; high-fidelity simulations being integral to launch risk analysis and component redesign.

For future flights, debris analysis has been directed to an assortment of shapes (typically thin and conical) that can potentially be shed from the external tank (ET) foam. The debris sources and their aerodynamic characteristics are put into the debris transport code, which calculates trajectory information to assess the potential damage or risk by a specific debris source to a specific structural component, such as the ET foam impacting the orbiter wing. For this interactive process to be effective, the debris transport analysis must be done rapidly. A single trajectory calculation requires 30–60 CPU-hours on the Columbia system, which provides enough throughput to rapidly and efficiently run hundreds of trajectories in a day, using only a fraction of the computational resources. A system with the power of Columbia is absolutely essential to run the typically hundreds of thousands of trajectories analyzed over the entire vehicle for each iteration. Results in Fig. 1 show that the average drag for the oscillating trajectory of an idealized frustum and the tumbling trajectory of a highly asymmetric debris piece are similar. Also note that farther downstream the debris travels before impact, the greater is the impact kinetic energy, as the aerodynamic drag is constantly increasing the relative velocity between the debris and the orbiter.

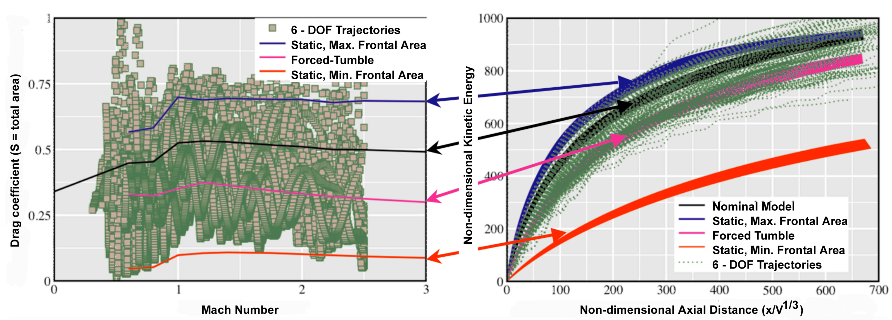

Fig. 1. Comparison of drag (left) and kinetic energy (right) for various debris shapes released at Mach 2.5. Symbols show unconstrained 6-DOF trajectories compared with trajectories using a nominal drag model based on the ensemble-averages.

This is not the case when considering the crossrange behavior (see Fig. 2). The dynamically stable oscillating frustum generates almost no crossrange, as the lift force oscillates first in one direction and then in the other, with little net effect. In order to provide a representative distribution, researchers used several shapes to develop the crossrange constraints. These include real digitized shapes, idealized frustums, ideal frustums with the center of mass offset, and slightly asymmetric shapes such as elliptical frustums with the small diameter slightly offset from the larger. The crossrange envelopes show a zero-lift trajectory emanating from the ET flange region. The intersection of this cone with the SSLV indicates that the fuselage and wing of the orbiter have potential for debris impacts from this flange location, along with regions of the left solid rocket booster.

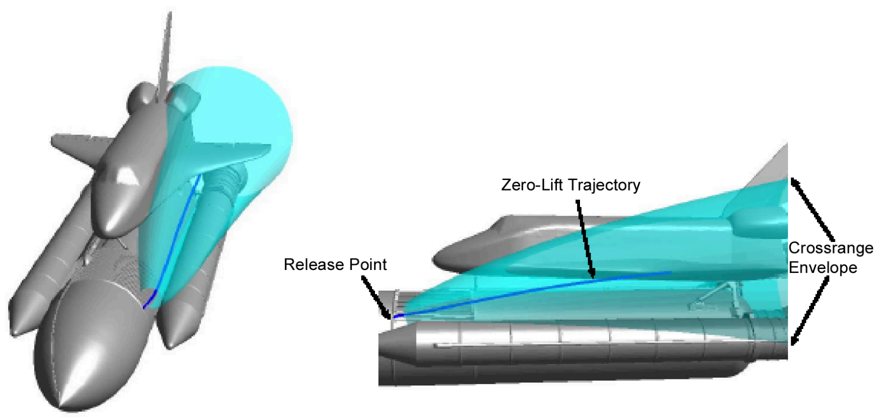

Fig. 2. Crossrange envelope superimposed upon the computed ballistic zero-lift trajectory. A statistical distribution of the crossrange within the envelope can be used for risk analysis.

The emphasis of this work is the development of an efficient process for modeling debris beyond ET insulating foam, including the ET liquid oxygen (LOX) frost ramps, insulating cork on the solid-rocket boosters, frost and ice on the ET acreage regions, and ice that can form on the ET feedline brackets. The flexibility of the modeling and simulation capability and the computing resources provided by Columbia allows the dynamic behavior of these diverse debris sources to be analyzed in a systematic and timely manner [16].

2.2 Space Shuttle Main Engine Flowliner Analysis

In May 2002, numerous cracks were found in the SSME #1 flowliner; specifically, at the gimbal joint in the liquid hydrogen (LH_2) feedline flowliner. Since then, repairs have been made to existing cracks on all orbiters. Long-term scientific investigations continue, because the root cause of the original cracks was not conclusively established and remaining Shuttle flights are involved until 2010.

High-fidelity computations have been conducted on the Columbia supercomputer to analyze the SSME LH_2 feedline flowliner [12]. Numerous computational models were used to characterize the unsteady flow features in the turbopump, including the Low-Pressure-Fuel-Turbopump (LPFTP) inducer, the orbiter manifold, and an experimental hot fire test article representing the manifold. Findings show that unsteady flow stemming from the LPFTP inducer is one of the major contributors to high-frequency cyclic loading that results in fatigue damage to the flowliners.

The flow fields for the orbiter manifold and the hot fire test article were computed and analyzed on Columbia using the INS3D incompressible Navier-Stokes flow solver [10, 11, 13] The first computational model included only the LPFTP inducer; by studying it, scientists were able to compare unsteady pressure values against existing data. To resolve the complex geometry in relative motion, an overset grid methodology [8] was employed, containing 57 overlapping zones with 26.1 million grid points. The second computational grid system, consisting of 264 overset grids with 65.9 million grid points, added the flowliner geometry and is shown in Fig. 3. The flowliner component alone contained 212 grids and 41 million points.

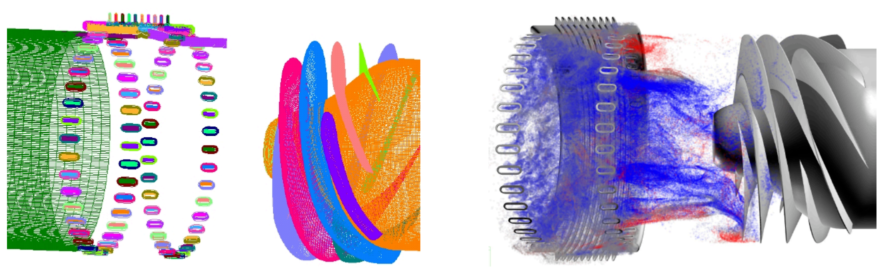

Fig. 3. Computational model for LPFTP inducer and the LH_2 flowliner: Grid (left), and computed results illustrating unsteady interaction of the flow in the bellows cavity and the back flow from the inducer (right).

To accelerate the grid generation process, scripts were developed to automatically create grids for each type of component. The size of the simulation is large, requiring parallel processing to obtain solutions with reasonable turnaround times. Two parallel programming paradigms were leveraged in the INS3D code: Multi-Level Parallelism (MLP) [23] and hybrid MPI+OpenMP. Performance for the two programming models was similar; however, only the MPI+OpenMP implementation can be executed on multiple nodes. Multi-node computations on Columbia showed that point-to-point

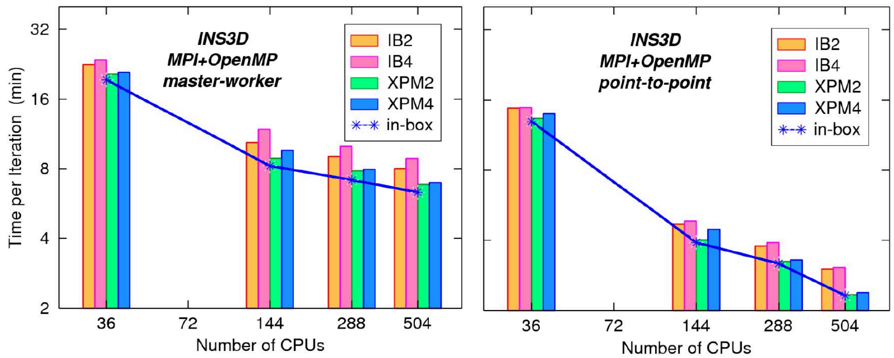

Fig. 4. INS3D performance on multiple nodes of Columbia.

implementation of the MPI communication performs more efficiently than the master-worker version [5] (see Fig. 4).

Results of the CFD calculations confirmed the presence of back flow caused by the LPFTP inducer. The region of reverse flow extended far enough upstream to interfere with both flowliners in the gimbal joint. Computed results for the test article were verified by correlation with pressure measurements, and confirmed a strong unsteady interaction between this back flow and the secondary flow in the bellows cavity through the flowliner slots. It was observed that a swirl on the duct side of the downstream flowliner is stronger than on the same side of the upstream flowliner, causing significantly stronger unsteady interactions through the downstream slots. This turbopump application using INS3D currently exhibits some of the best scalability performance on the Columbia system.

2.3 Crew Exploration Vehicle Abort Risk Assessment

Researchers are running high-fidelity CFD codes on Columbia to guide the designs of future space vehicles, including the CEV, and are building realistic models and developing new technologies to simulate flight risks for these new spacecraft. Risks and performance issues during both the ascent and entry-descent-landing phases are being carefully analyzed. The CEV will replace the Shuttle in 2010, and transport a maximum of six crew members to and from the International Space Station and up to four astronauts to and from the moon.

The CEV design includes a Launch Abort System (LAS) for crew escape, similar to that used in the Apollo capsule. Several computational modeling and simulation tools for analyzing abort scenarios have recently been developed and enhanced for use on Columbia. Under the simulation assisted risk assessment (SARA) project, NASA researchers have developed a probabilistic risk assessment (PRA) approach and demonstrated how risk analysis can be applied to launch abort using the Apollo configuration [14]. A PRA identifies the best level of fidelity for modeling critical failure modes associated with launch abort. Columbia is then used to conduct higher-fidelity modeling on specific failure modes. Two failure modes examined so far include booster explosion and those caused by re-contact with the booster during separation.

Analysis of the booster failure mode using Apollo data showed a possible catastrophic failure, leading to detonation of the propellant and creating blast wave overpressures that could fatally damage the LAS (see Fig. 5). As the risk model was being developed, it became apparent that the booster type and the nature of the failure it was likely to encounter, determined the environments

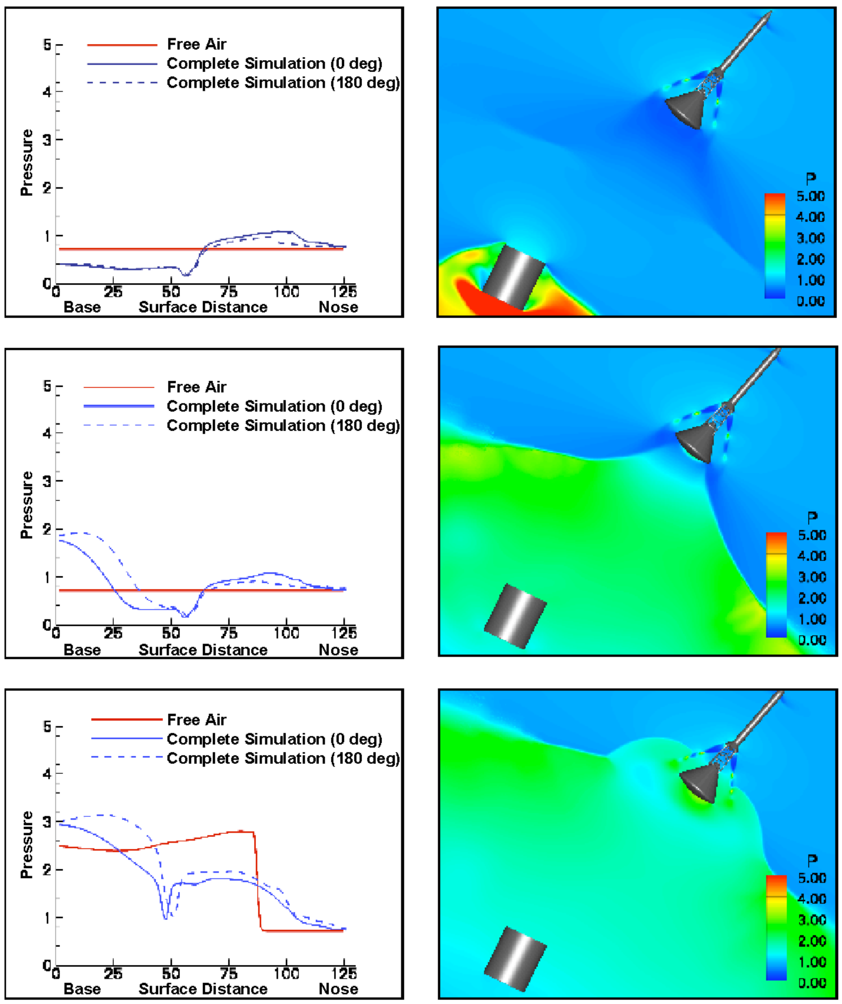

Fig. 5. Surface pressures (left) and flowfield (right) for blast wave propagating through wake of maneuvering LAS at $t = 41.6$, 72.0, and 85.5 msec (top to bottom).

under which the LAS must operate to ensure a successful abort. The process for characterizing this interaction must be carefully modeled and simulated.

One particular weakness found in an engineering-level model was the effect of headwind as the CEV ascends. To account for these effects in the risk analysis, high-fidelity blast wave models were built and simulated on Columbia using the Overflow Navier-Stokes code [15]. Results showed that headwinds significantly affect the nature and magnitude of the shock wave as it impacts an escaping CEV. This means that the warning time required to initiate the abort sequence is also affected. Additional work in high-fidelity simulations is being done to help engineers generate requirements for the LAS while taking headwind into consideration.

Another failure mode dependent on high-fidelity simulation involves the ability of the LAS to achieve clean separation of the CEV from the booster stack in the event of impending catastrophic failure. Simply put, the CEV must not scrape or re-contact the booster stack. This failure mode is particularly demanding because it involves complex proximity aerodynamics—modeling transonic flow as well as the complex flow at the small gap between the CEV and the booster stack at separation. Both Navier-Stokes simulations using Overflow, and Euler simulations using FlowCart [1], were conducted, and their results validated against transonic wind tunnel and abort flight test data from the Apollo era [4].

All these cases are computationally expensive to simulate. The complexity of the geometry and the flow-field required about 30 million grid points, which enabled good scalable performance up to 250 processors. About 20 cases were computed using Overflow at various ascent trajectories and separation thrust levels. Each case required approximately 20,000 CPU-hours on Columbia, including the computation of the initial steady-state solution. All failure modes benefited immensely from the HEC resources at the NAS facility. These tools and processes will likely be applied to analyze the actual LAS design, and to further understand the CEV failure modes and their impact on the vehicle's survivability.

3 Climate and Weather Applications

3.1 Global Ocean Modeling

To increase understanding and predictive capability for the ocean's role in future climate change scenarios, NASA has initiated a project called Estimating the Circulation and Climate of the Ocean, Phase II (ECCO2): High-resolution global-ocean and sea-ice data synthesis [17]. The goal is to produce increasingly accurate combinations of all available global-scale ocean and sea-ice data at resolutions that begin to resolve ocean eddies and other narrow current systems, which transport heat, carbon, and other properties in the ocean. These data syntheses are used to quantify the role of the oceans in the Earth's carbon cycle, understand recent changes in the polar oceans, and monitor time-evolving term balances within and between different components of the Earth system. This work aims to harness NASA's computational resources such as Columbia, advances in CFD and software engineering, and the ability to solve massive control problems.

The most challenging numerical experiment undertaken to date is a near-global simulation with $1/16\,^\circ$ horizontal grid spacing (approximately 6 km at the Equator and 1 km at high latitudes). The number of surface grid cells is about 25 million and the configuration has 50 vertical levels, bringing the total number of cells to just over 1.25 billion. Each of the 3D fields that describe the simulation domain and its time-evolving state requires 10 GB of storage. This configuration has been integrated on the 2,048-CPU sub-cluster of Columbia [9] (see Sec. 4). This workload is distributed evenly over 1,920 processors, so that each CPU is responsible for simulating about 586,000 grid cells (equivalent to a surface region roughly 210×210 km^2). Decomposing the workload over this many processors yields a setup that, with extensive diagnostics and analysis options included, uses about 870 MB of main memory per processor. With a timestep of two minutes, this performance allows a year of simulation to be completed in less than ten days.

To investigate solution convergence as horizontal resolution is increased, ECCO2 researchers are conducting a series of numerical simulations at $1/4\,^\circ$, $1/8\,^\circ$, and $1/16\,^\circ$ resolutions. Figure 6 shows significant changes in solution with varying resolution. Each plot captures the change in simulated sea-surface height due to eddy activity over a single month. Variation with resolution occur in regions where eddies are prevalent (such as the Gulf Stream, the Kuroshio, the Agulhas, the Drake Passage, and the Antarctic Circumpolar Current). For example, in the Gulf Stream, the area where the sea-surface height changes vigorously increases with higher resolution. Key behaviors, such as how tightly waters stick to the coast, or how far energetic eddies penetrate the ocean interior, also change significantly between resolutions.

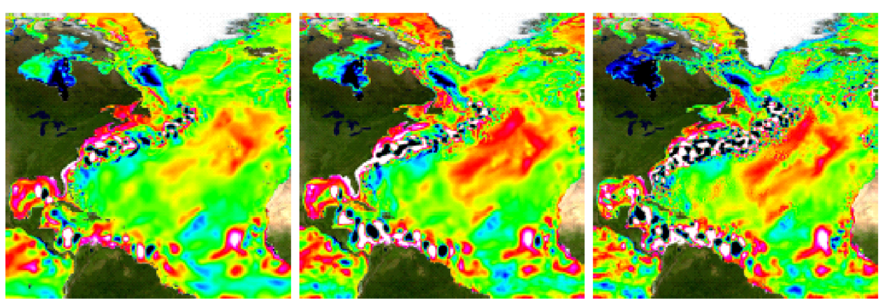

Fig. 6. Gulf Stream region sea-surface height difference plots for one month at resolutions of $1/4\,^\circ$, $1/8\,^\circ$, and $1/16\,^\circ$ (left to right). Color scale -0.125m to 0.125m.

Performance on Columbia shows that it is well suited for addressing these questions. The ECCO2 code achieves about 722 MFlop/s per CPU when running on 1,920 processors; this is 14 percent of the per-CPU Linpack benchmark performance achieved on Columbia [24]. The code consists of predominantly BLAS1 operations and cannot exploit the level of cache reuse that Linpack achieves. The scaling across multiple nodes is encouraging and suggests that configurations spanning eight or more 512-processor Altix systems—that would therefore support $1/20^\circ$ and higher resolutions—are within reach.

3.2 Atmospheric Modeling and Hurricane Prediction

The NASA Finite Volume General Circulation Model (fvGCM) is a unified numerical weather prediction (NWP) and climate model that could run on daily, monthly, decadal, and century time-scales. The model was originally designed for climate studies at a coarse resolution of about 2.5°, but has been running at much finer resolution on the Columbia supercomputer to answer the following question for NASA's mission in hurricane research [19]: How can weather/hurricane forecasts be improved and made more reliable over longer periods of time using computer modeling?

Hurricane forecasts pose challenges for general circulation models (GCMs), the most important being horizontal grid spacing. With the unique computing resources of Columbia, the model horizontal resolution was rapidly increased to $1/4^\circ$ in 2004 [2], and $1/8^\circ$ in early 2005. Recently, researchers have tested a $1/12^\circ$ resolution version, which is the first global weather model with a single-digit resolution (9 km at the Equator). A five-day forecast of total precipitable water with the $1/12^\circ$ degree fvGCM (see Fig. 7) shows fine scale weather events in the tropical area, which brings researchers to overcoming the fundamental barrier between global and mesoscale models [20].

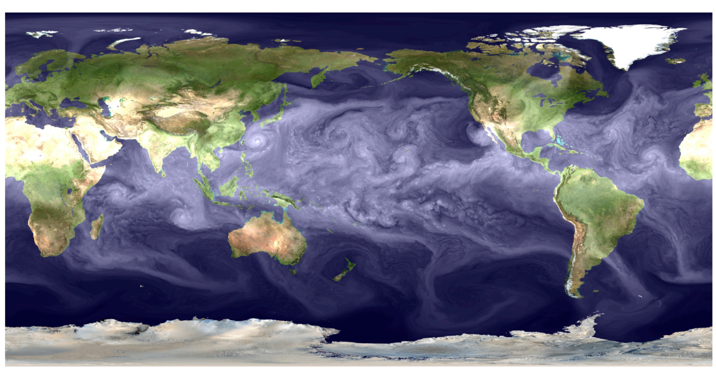

Fig. 7. Five-day forecasts of total precipitable water initialized in September 2004 with the $1/12^o$ fvGCM, giving a grid spacing of 9 km at the Equator.

During the 2004 hurricane season, the $1/4°$ model, which doubled the resolution adopted by most global models in operational NWP centers at that time, was running in real-time and providing remarkable landfall predictions up to five days in advance for major hurricane such as Charley, Frances, Ivan, and Jeanne [2]. Moreover, the model proved capable of resolving problems such as erratic track, abrupt recurvature, and intense extratropical transition. In the 2005 hurricane season, new research focused on the validation of the $1/8°$ fvGCM performance on hurricane forecasts, while the real-time $1/4°$ forecasts provided a baseline for comparisons. Being a global mesoscale-resolving model, the $1/8°$ resolution was the first to simulate mesoscale vortices (such as the Catalina Eddy and the Hawaiian Lee Vortex), which were generated by the interaction of the large-scale flows with better resolved surface forcing.

The 2005 Atlantic hurricane season was the most active in recorded history. There were 28 tropical storms and 15 hurricanes, four of which were Category 5. Accurate forecasts of these storms was a significant challenge to global and mesoscale modelers. It is well known that GCMs' insufficient resolutions undermine intensity predictions. Using the power of Columbia, NASA researchers demonstrated that this limitation could be overcome by performing six 5-day forecasts of hurricane Katrina [21] with the $1/8°$ fvGCM, showing promising intensity forecasts with small errors in center pressure of only ± 12 hPa. Notable improvement in Katrina's intensity forecasts occurred when the grid spacing decreased from $1/4°$ to $1/8°$, at which the near-eye wind distribution and the radius of maximum wind could be resolved more realistically. While the mesoscale-resolving fvGCM has produced very promising results for the past two years, significant potential for further advancement is still ahead.

4 Columbia Description

When the Columbia supercomputer became fully operational in October 2004, it increased NASA's computing capability ten-fold and helped revitalize the HEC effort within the U.S. Constructed in just four months, it has enabled scientists and engineers to perform important, breakthrough simulations in several arenas. Performing at a peak speed of 62 TFlop/s, Columbia has demonstrated its capability to support and accelerate the space agency's key missions and vision for exploration [6, 7].

Columbia is a 10,240-processor SGI Altix constellation composed of twenty 512-CPU nodes, twelve of which are model 3700, and the remaining eight are double-density 3700Bx2. Each node is a single-system-image (SSI) system, with 2 GB of shared memory per processor (1 TB per node). It uses SGI's NUMAflex global shared-memory architecture that allows access to all data directly and efficiently, without having to move them through I/O or networking bottlenecks. Within each node, the processors are interconnected via SGI's proprietary NUMAlink fabric. The 3700 utilize NUMAlink3 with a peak bandwidth of 3.2 GB/s, while the Bx2 have NUMAlink4 where the bandwidth is doubled to 6.4 GB/s.

The 20 Columbia nodes are connected by Voltaire InfiniBand fabric, as well as via 10- and 1-gigabit Ethernet connections. Four of the Bx2 nodes are tightly linked with NUMAlink4 (as well as the other fabrics) to form a 2,048-processor shared memory environment. Each processor in the 2,048-CPU subsystem is a 64-bit Intel Itanium2, running at 1.6 GHz, with 9 MB of level 3 cache and a peak performance of 6.4 GFlop/s. One other Bx2 node is equipped with these same processors. The remaining 15 nodes have Itanium2 processors running at 1.5 GHz, with 6 MB of level 3 cache, and a peak performance of 6.0 GFlop/s. Columbia is attached to more than 1 PB of online RAID storage through a Fibre Channel switch.

Each 512-processor Altix node has several salient features that make it particularly well-suited for executing large-scale compute and data-intensive applications that are interesting to NASA. For example, its less than 1 microsecond latency to memory significantly reduces the communication overhead. Typical problems are physics-based simulations involving a discretized grid of the physical domain that is partitioned across multiple processors exchanging boundary data at every time step. Columbia also was the first system (in November 2004) to exceed 1 TB/s memory bisection bandwidth on the STREAM benchmark [22]. With global shared memory and cache coherency, the nodes enable application programmers to use simpler and more efficient programming models than message passing. Problems requiring dynamic load balancing and/or adaptive gridding are much easier to program, control, and solve on Columbia, leveraging shared memory paradigms such as OpenMP and MLP [23]. In addition, the large shared memory of 1 TB per node allows bigger problems to remain resident on the system.

The development and operating environment on Columbia features a 64-processor SGI Altix front end, a Linux-based operating system, Altair PBS Professional job scheduler, Intel Fortran/C/C++ compiler, and SGI ProPack software.

5 Concluding Remarks and Future Vision

Over the last few decades, simulation methodologies have generally advanced along with computational technologies. Advanced tools have been developed to the point that many daily engineering and science problems can now be routinely computed; however, this is still done mostly using geometrically and or physically simplified or truncated models. Some of the physical models, such as those for turbulence and transition, and for high-temperature real gas, have not been advanced much more than what was available in the 1970s or early '80s.

To realize the full benefit of HEC, more inclusive modeling of geometry and physics is needed. Attempts to solve these problems have been made with some qualitative success. However, predictive capability is still very limited and prediction with accurate physics is yet to be accomplished; this will require inclusion of not only fluid dynamic quantities but other factors such as thermal loading, structural properties, and control. These computations will require not only larger computing resources but also increased data storage and sophisticated management technologies.

Many of Columbia's scientific and engineering users have stated that the system has allowed them to successfully complete investigations they never allowed themselves to dream of previously. Now, these users are envisioning what they can accomplish when even more powerful computing systems are available. NASA and the HEC community are working on developing petaflops-scale computers that can execute at rates more than 10^{15} operations per second. For example, the DARPA High Productivity Computer Systems (HPCS) program is working with several US computer vendors to provide a new generation of economically viable supercomputer systems before the end of this decade. The National Science Foundation (NSF) has a plan to provide computational resources to the general scientific community that can sustain petaflop performance by 2010. NASA is also planning its next-generation supercomputer system to meet the ever-increasing demand for computational resources required for a wide range of Agency-specific scientific discoveries and engineering applications.

For example, with NASA's next-generation system, scientists envision a launch simulation model designed to treat the entire launch environment until the vehicle has cleared the launch tower. The model would integrate 6-DOF multiple-body motion, debris transport and impact, propulsion system vibration and exhaust, acoustics due to exhaust, fuel accumulation in the exhaust plume, exhaust chemistry including fuel burning, thermal stress on the vehicle structure, and weather at the launch site. This very complex digital launch model would integrate data from propulsion simulation, meso-scale weather prediction, and experiment. Utilizing state-of-the-art flow simulation tools and petaflops-scale computing systems, researchers can attempt to compute a complete high-fidelity aerodynamic simulation with a realistic turnaround time—within a few days rather than several weeks.

In aerospace design, the most productive aspect of HEC applications has been to predict relative change among design variations. To push the limit of operation and to try bold new ideas, more predictive capabilities will be needed for complicated physical phenomena. Without accurate prediction, the capability impacts of HEC can be limited to the current level, even if more advanced facilities become available. To make these advances, high-fidelity computations using HEC facilities will be absolutely critical despite all the excitement about inexpensive PC clusters and distributed grid computing.

In Earth sciences, as resolutions increase significantly, the horizontal scales become smaller and the hydrostatic assumption is no longer valid. Non-hydrostatic dynamics, including eddy-resolving oceans, cloud-resolving atmosphere, and land models coupled with chemical and biological components, must therefore be included when running climate and weather applications. New schemes are also needed to better represent physical processes at less than 10-km resolutions. New grid systems (e.g. geodesic or icosahedral grids) are required due to inefficiencies of non-uniform latitude-longitude grids at ultra-high resolutions or convergence/stability problems at the poles.

A final note is related to human resources. Although modeling and simulation technology has advanced remarkably, many challenging cases require experts in computational physics. Computer science can automate a good portion of the simulation processes, thus saving a large amount of the human effort. However, blind application of tools without understanding capabilities and limitations of the methods involved could lead to catastrophic engineering results. As in many other engineering and science disciplines, modeling and simulation researchers and practitioners need to understand physics and the engineering systems being simulated. Experts who are willing to think through the flow physics in addition to software engineering, must still be developed for future generations.

References

1. M.J. Aftosmis, M.J. Berger, and G. Adomavicius, A parallel multilevel method for adaptively refined Cartesian grids with embedded boundaries, *38th AIAA Aerospace Sciences Mtg.*, AIAA Paper 2000-0808 (Reno, NV, 2000).
2. R. Atlas, O. Reale, B.-W. Shen, S.-J. Lin, J.-D. Chern, W. Putman, T. Lee, K.-S. Yeh, M. Bosilovich, and J. Radakovich, Hurricane forecasting with the high-resolution NASA finite volume general circulation model, *Geophys. Res. Lett.* 32(3), 2005, L03807, doi:10.1029/2004GL021513.
3. F.R. Bailey, High-end computing challenges in aerospace design and engineering, *3rd Intl. Conf. on CFD* (Toronto, Canada, 2004), pp. 13–26.
4. B.L. Barrier and O.C. Pendergraft Jr., Transonic aerodynamic characteristics of a powered wind-tunnel model of the Apollo Launch Escape Vehicle during separation, *NASA TM-X* 1336, 1967.
5. R. Biswas, M.J. Djomehri, R. Hood, H. Jin, C. Kiris, and S. Saini, An application-based performance characterization of the Columbia supercluster, *SC05* (Seattle, WA, 2005).
6. R. Biswas, D. Kwak, C. Kiris, and S. Lawrence, Impact of the Columbia supercomputer on NASA space and exploration missions, *2nd Intl. Conf. on Space Mission Challenges for Information Technology* (Pasadena, CA, 2006), pp. 51–58.
7. W. Brooks, M.J. Aftosmis, B. Biegel, R. Biswas, R. Ciotti, K. Freeman, C. Henze, T. Hinke, H. Jin, and W. Thigpen, Impact of the Columbia supercomputer on NASA science and engineering applications, *7th Intl. Wkshp. on*

Distributed Computing (Kharagpur, India, 2005), Springer LNCS 3741, 2005, pp. 293-305.
8. P.G. Buning, D.C. Jespersen, T.H. Pulliam, W.M. Chan, J.P. Slotnick, S.E. Krist, and K.J. Renze, Overflow user's manual—version 1.8g, *Technical Report*, NASA Langley Research Center (Hampton, VA, 1999).
9. C. Hill, D. Menemenlis, R. Ciotti, and C. Henze, Investigating solution convergence in a global ocean model using a 2048-processor cluster of distributed shared memory machines, *J. Scientific Programming*, submitted.
10. C. Kiris and D. Kwak, Numerical solution of incompressible Navier-Stokes equations using a fractional-step approach, *Computers & Fluids* 30(7), 2001, pp. 829–851.
11. C. Kiris, D. Kwak, and W.M. Chan, Parallel unsteady turbopump simulations for liquid rocket engines, *SC2000* (Dallas, TX, 2000).
12. C. Kiris, D. Kwak, W.M. Chan, and J. Housman, High-fidelity simulations for unsteady flow through turbopumps and flowliners, *44th AIAA Aerospace Sciences Mtg.*, AIAA Paper 2006-0089 (Reno, NV, 2006).
13. D. Kwak, J.L. Chang, S.P. Shanks, and S. Chakravarthy, An incompressible Navier-Stokes flow solver in three-dimensional curvilinear coordinate systems using primitive variables, *AIAA J.* 24(3), 1996, pp. 390–396.
14. S. Lawrence, D.L. Mathias, G. Klopfer, S. Pandya, M.E. Olsen, J. Onufer, T. Holst, and K. Gee, Simulation assisted risk assessment, *44th AIAA Aerospace Sciences Mtg.*, AIAA Paper 2006-0090 (Reno, NV, 2006).
15. S. Lawrence, D.L. Mathias, K. Gee, and M.E. Olsen, Simulation assisted risk assessment: Blast overpressure modeling, *8th Intl. Conf. on Probabilistic Safety Assessment and Management*, Paper 0197 (New Orleans, LA, 2006).
16. D.J. Mavriplis, M.J. Aftosmis, and M.J. Berger, High resolution aerospace applications using the NASA Columbia supercomputer, *SC05* (Seattle, WA, 2005).
17. D. Menemenlis, C. Hill, A. Adcroft, J.-M. Campin, B. Cheng, R. Ciotti, I. Fukimori, P. Heimbach, C. Henze, A. Köhl, T. Lee, D. Stammer, J. Taft, and J. Zhang, NASA supercomputer improves prospects for ocean climate research, *Eos Trans. AGU* 86(9), 2005, pp. 89,95–96.
18. S.M. Murman, M.J. Aftosmis, and S.E. Rogers, Characterization of Space Shuttle ascent debris aerodynamics using CFD methods, *43rd AIAA Aerospace Sciences Mtg.*, AIAA Paper 2005-1223 (Reno, NV, 2005).
19. NASA's Role in Hurricane Research, http://www.nasa.gov/pdf/147432main_hurr_fact_sheet.pdf.
20. B.-W. Shen, R. Atlas, J.-D. Chern, O. Reale, S.-J. Lin, T. Lee, and J. Chang, The 0.125 degree finite-volume general circulation model on the NASA Columbia supercomputer: Preliminary simulations of mesoscale vortices, *Geophys. Res. Lett.* 33(5), 2006, L05801, doi:10.1029/2005GL024594.
21. B.-W. Shen, R. Atlas, O. Reale, S.-J. Lin, J.-D. Chern, J. Chang, C. Henze, and J.-L. Li, Hurricane forecasts with a global mesoscale-resolving model: Preliminary results with Hurricane Katrina (2005), *Geophys. Res. Lett.* 33(13), 2006, L13813, doi:10.1029/2006GL026143.
22. STREAM Benchmark, http://www.streambench.org/.
23. J.R. Taft, Achieving 60 Gflop/s on the production CFD code Overflow-MLP, *Parallel Computing* 27(4), 2001, pp. 521–536.
24. Top500 Supercomputer Sites, http://www.top500.org/.